Towards Automatic Learning of Heuristics for Mechanical Transformations of Procedural Code

Guillermo Vigueras Manuel Carro
IMDEA Software Institute
Campus de Montegancedo 28223
Pozuelo de Alarcón, Madrid, Spain
{guillermo.vigueras,manuel.carro}@imdea.org

Salvador Tamarit Julio Mariño
Universidad Politécnica de Madrid
Campus de Montegancedo 28660
Boadilla del Monte, Madrid, Spain
{salvador.tamarit,julio.marino}@upm.es

arXiv:1603.03022v2 [cs.PL] 10 Mar 2016

Abstract

The current trend in next-generation exascale systems goes towards integrating a wide range of specialized (co-)processors into traditional supercomputers. However, the integration of different specialized devices increases the degree of heterogeneity and the complexity in programming such type of systems. Due to the efficiency of heterogeneous systems in terms of Watt and FLOPS per surface unit, opening the access of heterogeneous platforms to a wider range of users is an important problem to be tackled. In order to bridge the gap between heterogeneous systems and programmers, in this paper we propose a machine learning -based approach to learn heuristics for defining transformation strategies of a program transformation system. Our approach proposes a novel combination of reinforcement learning and Classification methods to efficiently tackle the problems inherent to this type of systems. Preliminary results demonstrate the suitability of the approach for easing the programmability of heterogeneous systems.

Categories and Subject Descriptors I.2.6 [*Artificial Intelligence*]: Learning; C.1.4 [*Processor Architectures*]: Parallel Architectures

General Terms Learning, Parallelism, Optimization

Keywords Program Transformation, Machine Learning, Heterogeneous Systems

1. Introduction

The introduction of multi-core processors marked the end to an era dominated by a constant increase of clock-rate, following Moore's law. Today, even the strategy to increase performance through higher core-counts is facing physical obstacles and, more importantly, power consumption issues which are currently one of the main limitations for the transition from petascale systems to next-generation exascale systems. Hence, the metric for measuring performance improvement of computational platforms is no longer absolute FLOPS numbers but rather FLOPS per Watt and FLOPS per surface unit of hardware.

Thus, rather than having CPU-based supercomputers a growing trend goes towards integrating a wide range of specialized (co-)processors into traditional supercomputers. These specialized architectures can outperform general purpose ones while requiring lower energy consumption values and also less hardware real state. A study of energy consumption in data centers (Koomey 2008), has revealed the impact that efficient computer platforms can have in the worldwide energy consumption. This study analyzed historical energy consumption data for the time period 2000-2005 and 2005-2010. The study results revealed that in 2005 data centers used 0.97% of the worlds total electricity consumption and in 2010 the value increased up to 1.3%. However, the predicted value for 2010 based on the trends in the time period 2000-2005 would have been 2.2% instead of 1.3%, as reported. The study explained that the reduction was mainly a consequence of an increase in the use of energy-efficient computational platforms, showing the economical impact of using *greener* architectures.

However, the integration of different specialized devices increases the degree of heterogeneity of the system. This high degree of heterogeneity makes programmability of such systems only possible for a few experts, and almost prohibits portability of the application onto different resource infrastructures. Classical programming models are all designed for (a) Turing-inspired sequential execution and (b) von Neumann-like memory architecture, whereas a compiler typically optimizes code only for one (homogeneous) destination infrastructure at a time. Heterogeneous destinations typically require all programming tasks for hybrid architectures to be manually done by the developer. In fact, no common current programming model manages to ease the programmability of heterogenous platforms and the portability of applications to this type of platforms while keeping good performance.

It is crucial to note in this context, that programming models are only likely to be taken up by the programming communities if they (a) are coherent with classical approaches, (b) are easy and intuitive and (c) show clear performance / usability improvements. Accordingly, development of new programming models that hide low-level details and ease the programmability of the range of architectures available in heterogenous systems is slow, even if some programming models have already been proposed in the past for each individual architecture (e.g. OpenCL (Stone et al. 2010)).

Opening the access of heterogeneous platforms to a wider spectrum of users is an important issue to be tackled. In past decades, different scientific areas have shown an enormous advance and development thanks to the use of massive parallel architectures to implement *in-silico* models. The use of "virtual labs" through computational platforms has allowed to better understand the physical phenomena under study, investigate a much wider range of solutions and drastically reduce the cost with respect to performing real experiments. Since the exascale computing era is pushed and driven by the computational requirements of scientific and industrial applications, new programming models, tools and compilers need to be released in order to open the access of exascale systems to a wider programming community and allow researchers from different communities to apply their findings to the society and improve human welfare.

In order to ease the programmability of heterogeneous platforms, we propose in this paper a machine learning -based approach to learn heuristics to guide the code transformation process of a rule-based program transformation system (Tamarit et al. 2015).

This type of systems pose different problems like the search-space exploration problem arising from the application of transformation rules in arbitrary orders or the definition of a stop criteria for the transformation system. For the latter we propose the use of classification trees and for the former we propose a novel approach based on reinforcement learning . We have performed a preliminary evaluation of the approach obtaining promising results that demonstrate the suitability of the approach for this type of transformation systems.

The rest of the paper is organized as follows. Section 2 reviews previous work in the field of program transformation systems in general and previous approaches using machine learning techniques. Section 3 describes the toolchain where our machine learning -based approach is integrated. Section 4 describes the abstraction of code defined in order to apply machine learning methods. Section 5 describes the methods used to learn program transformation strategies. Later, Section 6 presents some preliminary results and finally Section 7 summarizes the conclusions and proposes future work.

2. State of the Art

Rule-based program transformation systems support the formulation of basic code transformation steps as generic rules and arrange their automatic application. This scheme offers the flexibility of splitting complex code transformations in small steps, admitting efficient implementations that scale to large programs. By adding more rules, the transformation system can also increase its capabilities in a scalable way. Rule-based systems also allow to decouple the definition of transformation steps (i.e. rule applications) from the strategies followed to apply the rules. This decoupling provides more flexibility to try different transformation strategies and select the best one according to the purpose of the system (Bagge et al. 2003; Schupp et al. 2002). Rule-based transformation has been used before to generate code for different computational platforms.

The transformation of C-like programs an its compilation into synchronous architectures, like FPGAs, and asynchronous platforms, like multi-core CPUs, have been addressed before (Brown et al. 2005). However, the input language of the approach (Handel-C) is meant to specify synchronous systems, thus limiting its applicability to this type of systems. A completely different approach is to use linear algebra to transform the mathematical specification of concrete scientific algorithms (Franchetti et al. 2006; Fabregat-Traver and Bientinesi 2013; Di Napoli et al. 2014). Here, the starting point is a mathematical formula and, once the formula is transformed, code is generated for the resulting expression. However, a reasonable acceleration over hand-tuned code happens only for those algorithms, and applying the ideas to other contexts does not seem straightforward.

Machine learning techniques have been already used in the field of compilation and program transformation (Mariani et al. 2014; Pekhimenko and Brown 2010; Agakov et al. 2006). All these approaches share the same principles as they obtain an abstract representation of the input programs in order to apply machine learning methods. Nevertheless, previous approaches target some specific architecture limiting the applicability of the approach and making it not suitable for heterogeneous platforms. In our approach, we obtain program abstractions but in order to enable the machine learning -based transformation and compilation of a program written in C for heterogeneous systems. Additionally, none of the previous works have explored the use of reinforcement learning methods (Kaelbling et al. 1996) in the field of program transformation and compilation.

3. Program Transformation Toolchain for Heterogeneous systems

In this work we propose the automatic learning of transformation strategies for a rule-based transformation toolchain (Tamarit et al. 2015, 2016). This type of systems offer a higher flexibility over compilation infrastructures where the code transformations and optimizations are embedded in the compilation engine (Stallman and GCC Developer Community 2009). Rule-based systems support the formulation of basic code transformation steps as generic rules and arrange their automatic application. This scheme offers the flexibility of splitting complex code transformations in small steps, admitting efficient implementations that scale to large programs. By adding more rules, the transformation system can also increase its capabilities in a scalable way. Rule-based systems also allow to decouple the definition of transformation steps (i.e. rules) from the strategies followed to apply the rules. This decoupling provides more flexibility to try different transformation strategies and select the best one according to the purpose of the system (Bagge et al. 2003; Schupp et al. 2002).

The rule-based transformation toolchain mentioned before (Tamarit et al. 2015, 2016) defines a transformation process consisting of two main stages, as shown in Figure 1. The first one is a transformation phase in which the input C code is transformed into semantically equivalent c Code but which is better suited for a given platform. It basically reshapes the input code taking into account syntactic/semantic restrictions of the compilers/programming models of each target platform (e.g. if the target compiler does not accept general `if` statements, the input code will be transformed to remove them if possible, without modifying the program semantics). In the first stage, transformation rules are formalized in an internal domain-specific language that takes care of the syntactic and semantic conditions required to apply a given transformation. Transformation rules operate at *abstract syntax tree* (AST) level transforming an input code that matches a pattern into another code, if rule conditions are met. In this way, this first phase defines a search space exploration problem where the different nodes of the search space would be the different states of the input code obtained as a result of applying code transformations and the transitions among the search space nodes would be the transformation rules applied. Thus, the high number of rules in the transformation system results in a combinatorial explosion of the code states to explore.

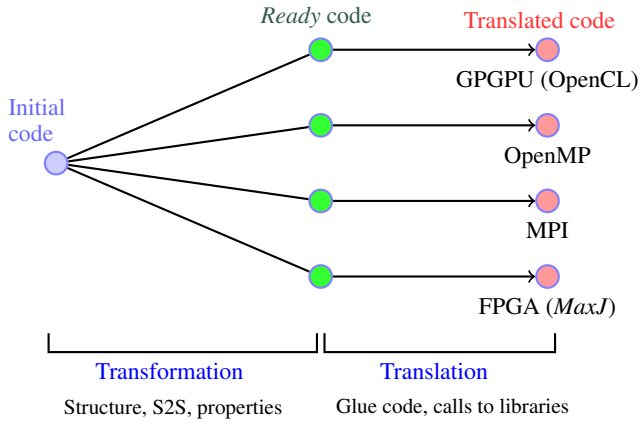

Figure 1. Stages of the transformation tool

The second phase of the toolchain in Figure 1 consists of a translation process which generates code for a given platform taking as

input the code generated in the first phase. Once platform-specific code is generated, it is compiled using a platform-specific compiler. This second phase poses another problem to be tackled: the definition of the starting condition for the translation phase, which is also the stop condition of the transformation process. Since depending on program complexity some platforms might require a long compilation process taking up to tens of hours, like for an FPGA, the toolchain should translate and compile only code that is likely to be successfully compiled on a given target platform rather than trying to compile every single code obtained during the transformation phase. Additionally, when considering the first and second phases together, the transformation sequences obtained in the first phase from the chaining of several rules, might not improve monotonically non-functional metrics associated to an input code (e.g. execution time, energy consumption, etc.), but still produce code with better performance in the second phase.

Given the nature of the problems associated with the transformation toolchain, we propose the use of different machine learning techniques to effectively guide the transformation process. On one hand, we have used reinforcement learning methods to learn transformation strategies that take into account the non-monotonic behavior of the transformation sequences. On the other hand, we have used classification methods to learn appropriate states that can be translated for the target platforms. Also machine learning techniques require descriptions of the problem domain. For that reason, we compute an abstract representation of the input programs and the intermediate codes generated during the transformation process.

4. Code-Abstraction Mapping

As mentioned in previous section, machine learning methods operate on descriptions of the problem domain. In our case, we need to obtain abstractions of procedural code written in C. These abstractions should reflect the code changes performed by the transformation rules which operate at the AST level. At the same time, mapping of code to abstractions must be univocal in order to avoid conflicts and ensure a correct behavior of machine learning methods. For that reason, we represent abstractions as quantitative descriptions that capture changes performed by rules on code. These quantitative descriptions involve features like: AST patterns, control flow, data layout, data dependencies, etc. To obtain these abstractions we have developed a static code analysis (SCA) tool that parses the AST to extract the features mentioned before. The SCA tool has been implemented in Python using the *pycparser*[1] module and can obtain feature information either by directly analysing the code or by parsing code annotations provided by the user or by third-party analysers. The set of code annotations considered in the abstractions coincides with the annotations supported by the transformation toolchain mentioned before (Tamarit et al. 2015).

Currently, code abstractions consist of a vector formed out of fifteen features. Some of them capture changes in code performed by transformation rules and also reflect code patterns that match the syntactic/semantic restrictions of target compiler/programming models:

- **Maximum nested loop depth:** This feature measures the depth of nested `for` loops. A value of 0 for this feature means no nested loops are present, a value of 1 means two nested loops, a value of 2 means three nested loops and so on.
- **Number of function calls:** This feature counts the number of function calls present in the analysed code.
- **Number of array writes shifted within for loops:** This feature counts the number of array accesses with positive offset for the following pattern: a `for` loop with writes accesses of multiple positions of the same array within the same loop iteration and some of the write accesses use a positive offset with respect to the loop iteration variable.

  The following two codes explain the meaning of this feature. For the code on the left, this feature would take a value of 1 since for each iteration there are multiple write accesses, one of them with positive offset (i.e. `i+1`). However, for the code on the right this feature takes a value 0 since for each iteration only position `i+1` is written even if it is done with positive offset. The multiple writes within the same loop iteration is detected by this feature since a code with this pattern will fail to be compiled for a FPGA.

```
for(i=1;i<N;i+=2) {
   v[i] = v[i-1];
   v[i+1] = v[i-1]*i;
}
```

```
for(i=0;i<N-1;i++) {
   aux = i*i;
   v[i+1] = aux;
}
```

- **Irregular loops:** This feature takes a binary value and informs whether loops contain (1) or not (0) statements like `break` or `continue` that alter the normal flow a program.
- **Global variables:** This feature takes a binary value and states if any global variable is written within the piece of code analysed (1) or not (0).
- **If statements:** This feature counts the number of `if` statements within the piece of code analysed.
- **Static limits of for loop:** This feature informs if all the `for` loops in the piece of code analysed do not have static iteration limits, i.e. if they do not change along the iteration space of the loop. The following code shows an example where the upper iteration limit of the inner loop is not static since it might change for each iteration of the outer loop if elements of array `v` are removed by function `clean`.

```
for(j=0;j<N;j++) {
   for(i=0;i<size(v);i++)
      update(v[i]);

   clean(v);
}
```

- **Iteration independent loop:** This feature counts the number of `for` loops in the code analysed without carried dependencies across iterations. This information is obtained from annotations provided by the user or by third-party code analysers.
- **Any for loop with loop_schedule:** This feature detects if there is some pattern composed of two nested loops used to iterate over an array split in chunks. This code pattern is detected through an annotation that can be inserted either by the user or by the transformation tool after applying a loop scheduling transformation. The following code shows an example of loop_schedule where array `v` is accessed in chunks of size N:

```
#pragma stml loop_schedule
for(j=0;j<M;j++) {
   w[j] = 0;
   for(i=0;i<N;i++)
      w[j] += v[j*N+i];
}
```

[1] https://github.com/eliben/pycparser

- **Number of loop invariant var:** This feature quantifies the number of variables that are assigned outside a `for` loop and are not modified within it.
- **Number of loop hoisted var modifications:** This feature is used to count the number of variables that are assigned outside a `for` loop and are modified within the loop.
- **Number of non-1D array:** This feature counts the number of arrays in the code with a number of dimensions higher than 1.
- **Number of auxiliary variables to access arrays:** This feature counts the number of auxiliary variables used to index an array. If a variable used as array index is not within the set of iteration variables of the `for` loops that iterates over the array, then it is considered as an auxiliary index variable. The following code shows a simple example where the `aux` variable is used to index the array `v` instead of using the iteration variable `j` of the `for` loop:

```
aux = 0;
for(j=0;j<N;j++) {
   w[i] = v[aux];
   aux++;
}
```

- **Total number of for loops:** This feature counts the total number of `for` loops independently of the nesting of loops.
- **Non-normalized for loops:** This feature counts the number of non-normalized `for` loops. Currently, a `for` loop is considered as normalized if the step of the iteration is 1 and it is considered as non-normalized otherwise.

The SCA tool in charge of obtaining the abstractions described above can be seen as a function mapping codes to abstractions, defined in the following way:

$$A: \ Code \rightarrow Abstraction$$

Function A is used later in Section 5.2 for explaining the mapping of codes to reinforcement learning states. Currently code abstractions consist of the fifteen features described before since they were sufficient to obtain some preliminary results for the set of use case applications described in Section 6. However, we plan to increase the vector of features as we increase the set of use case applications.

5. Automatic Learning of Transformation Heuristics

As mentioned before, the transformation engine that guides the toolchain must tackle some problems associated with rule-based code transformation systems. On one hand, the search space with non-monotonic behavior must be efficiently explored and on the other hand an effective stop criteria for the search procedure must be defined. Our approach uses reinforcement learning to solve the former problem and classification trees to solve the latter. Following subsections describe the details of these machine learning techniques. An example to show how it works in practice is provided later.

5.1 Classification Trees

In machine learning and statistics, classification is the problem of identifying the category to which a new observation belongs among a set of pre-defined categories. The classification is done on the basis of a training set of data containing observations for which it is known to which category they belong (Marsland 2009). Different formalisms can be used to classify. We have decided to start evaluating the adequacy of classification trees for our problem since it intuitively and implicitly performs feature selection without hard data preparation requirements.

A classification tree is a simple representation for classifying examples according to a set of input features. All of the input features have finite discrete domains, and there is a single target variable called the *classification* feature. Each element of the domain of the target variable is called a class. In a classification tree each internal (non-leaf) node is labeled with an input feature. Each leaf of the tree is labeled with a class or a probability distribution over the classes. Thus, a tree can be learned by splitting the source data set into subsets based on values of input features. This process is repeated on each derived subset in a recursive manner called recursive partitioning. The recursion is completed when the subset at a node has the same value of the target variable, or when splitting no longer improves the predictions. Typically, the source data comes in records of the form:

$$(\mathbf{x}, Y) = ([x_1, x_2, x_3, ..., x_k], Y)$$

The dependent variable, Y, is the target variable that the classification tree generalizes in order to be able to classify new observations. The vector $\mathbf{x}$ is composed of the input features x_i, used for the classification. In this context, the source input data for our problem is composed of the vectors of the abstractions described in Section 4, i.e., $k = 15$ in our case. The domain of the target variable can take values among the four different final platforms we currently support: FPGA, GPU, Shared-Memory CPU (SM-CPU) and Distributed-Memory CPU (DM-CPU). Since a given code and its associated abstraction might be well suited for more than one platform, we obtain 15 possible classes for our target variable. This number is computed as the sum all the combinations of four elements taken m at a time, where $1 \leq m \leq 4$. So the total number of classes N is computed as:

$$N = \sum_{m=1}^{4} C_4^m = 15, \ where \ C_n^m = \frac{n!}{m!(n-m)!}$$

The classes obtained for the target variable allow to define the final states of the transformation stage of the toolchain described in Section 3, which will also serve to define the final states for the reinforcement learning algorithm that is described in the following section. As a final remark, the classification-based learning described in this section have been implemented using the Python library *Scikit-learn* (Pedregosa et al. 2011). This library implements several machine learning algorithms and is widely adopted in the scientific community providing good support and ample documentation.

5.2 Reinforcement Learning

Reinforcement learning (Marsland 2009) is an area of machine learning concerned with how software agents ought to take actions in an environment so as to maximize some notion of cumulative reward. A reinforcement learning agent interacts with its environment in discrete time steps. At each time t, the agent receives an observation o_t, which typically includes the reward r_t. It then chooses an action a_t from the set of available actions, which is subsequently sent to the environment. The environment moves from current state s_t to a new state s_{t+1} providing the reward r_{t+1} associated with the transition (s_t, a_t, s_{t+1}). The goal of a reinforcement learning agent is to collect as much reward as possible.

According to the previous description, reinforcement learning seems well suited to represent the optimization process of a programmer or a compiler, since it typically consists of iteratively improving an initial program in discrete steps, where code changes correspond to actions and code versions obtained during the op-

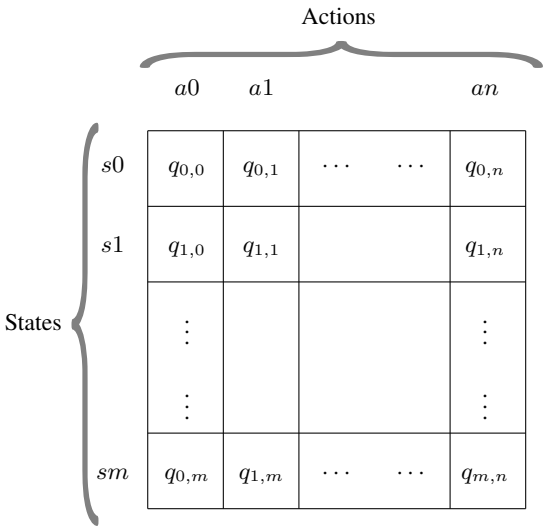

Figure 2. State-Action table Q. It will eventually be filled in with values $q_{i,j} \in \mathbb{R}$ obtained from the learning process.

timization process correspond to states. Moreover, code is typically evaluated after every change, often according to some non-functional properties such as execution time, memory consumption speedup factor, ... The result of these evaluations can be easily translated into rewards and penalties that support the learning process.

The result of the learning process of the agent is a *state-action* table (Figure 2) which will eventually contain values for each combination (s, a) of states and actions. These values are scores for the expected profit to be obtained from applying action a to state s. This table is initially filled in with a default value and is iteratively updated following a learning process which we briefly describe below.

The process of reinforcement learning is based on a set of predetermined transformation sequences, which are assumed to be models to learn from. Each sequence S is composed of a set of states $S = s_0, s_1, \ldots, s_{t-1}, s_t, s_{t+1}, \ldots, s_l$ and the actions which transform one state into the next one. The final state of each transformation sequence has a different reward value related, in our case, with the performance of the final code corresponding to state s_l in the sequence (better performance gives higher rewards). The reinforcement learning training phase consists of an iterative, stochastic process in which states (i.e. code abstractions) from the training set are randomly selected. For each state s a *learning episode* is started by selecting the action a with the highest value in Q for that s and moving to a new state s' according to transition (s, a, s'). From state s' the same transition process is repeated to advance in the learning episode until a final state is reached or a given number of episode steps is performed. When the episode terminates, the values in Q corresponding to the states and actions of the visited sequence are updated according to Equation 1, where $Q_{init}(s_t, a_t)$ is the initial value of Q for state s_t and action a_t. Note that s_t (resp. a_t) is not the t-th state in some fixed ordering of states, but the t-th state in the temporal ordering of states in the sequence used to learn.

$$Q(s_t, a_t) = \begin{cases} Q(s_t, a_t) + \alpha \cdot (r_{t+1} + \gamma \cdot Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t)) & \text{if } s_t \; not \; final \\ Q_{init}(s_t, a_t) & \text{otherwise} \end{cases} \quad (1)$$

The final states in Equation 1 are defined based on the classification as described in Section 5.1. Two additional parameters appear in Equation 1: the learning rate $\alpha, 0 < \alpha \leq 1$, and the discount factor $\gamma, 0 < \gamma \leq 1$. The learning rate determines to what extent the newly acquired information will override the old information. A factor of 0 will make the agent not learn anything, while a factor of 1 would make the agent consider only the most recent information. The discount factor implements the concept of *delayed reward* by determining the importance of future rewards. A factor of 0 will make the agent opportunistic by considering only current rewards, while a factor close to 1 will make it strive for a long-term high reward. If the discount factor reaches or exceeds 1, the values in Q may diverge (Marsland 2009).

In order to use the reinforcement learning *state-action* in our setting, we need to define some mappings. The abstraction of a concrete piece of code is provided by function A (Section 4). Abstractions and transformation rules must be mapped to states and actions, respectively, in order to index the *state-action* table. This mapping is done through functions SM and AM, defined as:

$$SM : Abstraction \rightarrow State$$

$$AM : Rule \rightarrow Action$$

Based on the mapping of abstractions and rules defined, the reinforcement learning *state-action* table of Equation 1 can also be modeled as a function Q:

$$Q : State \times Action \rightarrow \mathbb{R}$$

Using functions A, SM, AM and Q, the strategy of the transformation toolchain for selecting rules at each transformation step can be modeled with a function RS:

$$RS : Code \rightarrow Rule$$

This function takes as input a given code c and selects the transformation rule ru associated to action $AM(ru)$, that maximizes the value provided by Q for the state $SM(A(c))$ associated to input code c. Thus, the rule selection strategy can be expressed as:

$$RS(c) = \arg\max_{ru \in Rule} Q(SM(A(c)), AM(ru))$$

The operator arg max may return, by definition, the empty set, a singleton, or a set containing multiple elements. However, in our problem parameters α and γ as well as the reward values r_{t+1} appearing in Equation 1 can be tuned to ensure that a single rule is returned, this avoiding a non-deterministic RS function. Section 6 gives further details on how we selected their values.

The relationships among the code, its abstraction, the rules, and the contents of the state-action matrix are depicted in Figure 3. Table Q is used as follows: for a concrete code c_k we find its abstraction $\mathcal{C}_i = A(c_k)$. Let us assume $i = 0$. From the row i corresponding to $\mathcal{C}_i$ in matrix Q we obtain the column j with the highest value $q_{i,j}$ (in our example, $q_{0,1}$, in blue and boldface). Column j corresponds to rule R_j, which is expected to give the most promising code when applied to a code state whose abstraction is $\mathcal{C}_i$ (in our case it would be R_1). Rule R_j would be applied to c_k to give c'. If c' corresponds to a final state, the procedure finishes. Otherwise we repeat the procedure taking c' as input and finding again a rule to transform c'.

In order to implement the reinforcement learning -based module in charge of learning heuristics for program transformation we have used the Python library *PyBrain* (Schaul et al. 2010). This library adopts a modular structure separating in classes the different

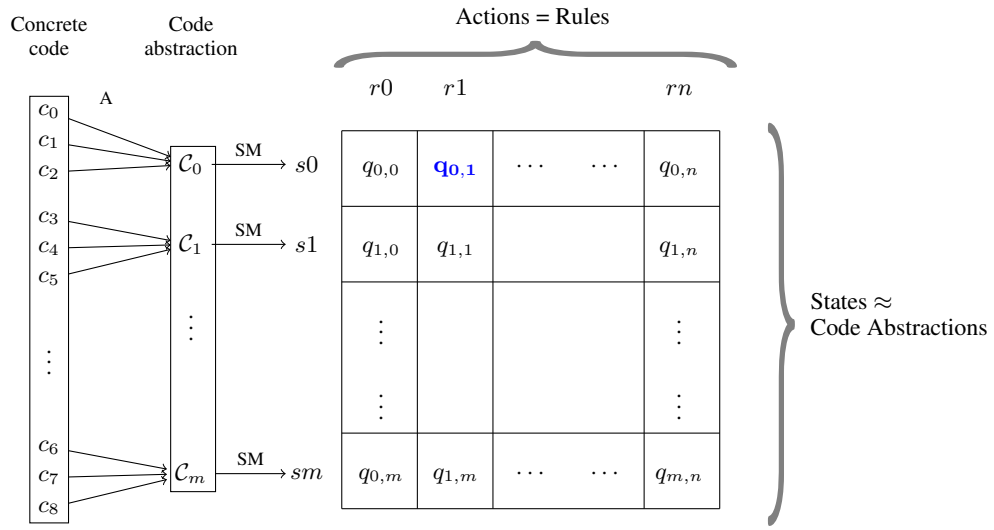

Figure 3. State-Action table for code, code abstraction, and rules.

concepts present in reinforcement learning like the environment, the observations and rewards, the actions, etc. This modularity allowed us to extend the different classes and ease their adaptation to our problem. The *PyBrain* also provides flexibility to configure the different parameters of the reinforcement learning algorithm.

5.3 Simple Example

This section uses a 2D convolution kernel as an example to show the resulting *state-action* table obtained after learning from a simple transformation sequence with five states and two transformation rules. The first rule (R_0) considered in the example transforms a non-1D array into a 1D array. The second rule (R_1) performs a collapse of two nested for loops producing a single loop. The initial, intermediate and final codes obtained from the application of these rules is described below. Listing 1 through Listing 5 highlights changes in code using this style for indicating the portion of the code that will be changed after rule application. Highlight code using this style indicates the resulting code after applying some transformation rule.

Listing 1 shows the initial code and the associated vector of features, as a code comment , according to the description in Section 4.

Listing 1. Initial Code

```
// ABSTR: [3, 0, 0, 0, 0, 0, 1, 0, 0, 1, 1, 3, 2, 4, 0]

int dead_rows = K / 2;
int dead_cols = K / 2;

int normal_factor = K * K;

for (r = 0; r < N - K + 1; r++) {
  for (c = 0; c < N - K + 1; c++) {
    sum = 0;
    for (i = 0; i < K; i++) {
      for (j = 0; j < K; j++) {
        sum += input_image[r+i][c+j] * kernel[i][j];
      }
    }
    output_image[r+dead_rows][c+dead_cols] = (sum /
        normal_factor);
  }
 }
```

Listing 2 shows the result of applying rule R_0 to the code in Listing 1. It can be seen that the array `input_image` is transformed into a one dimensional array and the vector of features associated to code changes accordingly.

Listing 2. Transformation step 1

```
// ABSTR: [3, 0, 0, 0, 0, 0, 1, 0, 0, 1, 1, 2, 2, 4, 0]

int dead_rows = K / 2;
int dead_cols = K / 2;

int normal_factor = K * K;

for (r = 0; r < N - K + 1; r++) {
  for (c = 0; c < N - K + 1; c++) {
    sum = 0;
    for (i = 0; i < K; i++) {
      for (j = 0; j < K; j++) {
        sum += input_image[(r+i)*(N - K + 1) +
            (c+j)] * kernel[i][j];
      }
    }
    output_image[r+dead_rows][c+dead_cols] = (sum /
        normal_factor);
  }
 }
```

Listing 3 shows the result of applying again rule R_0 to the code in Listing 2. It can be seen that now the array `kernel` is transformed into a one dimensional array and the vector of features associated to code changes accordingly.

Listing 3. Transformation step 2

```
// ABSTR: [3, 0, 0, 0, 0, 0, 1, 0, 0, 1, 1, 1, 2, 4, 0]

int dead_rows = K / 2;
int dead_cols = K / 2;

int normal_factor = K * K;

for (r = 0; r < N - K + 1; r++) {
  for (c = 0; c < N - K + 1; c++) {
    sum = 0;
```

```
    for (i = 0; i < K; i++) {
      for (j = 0; j < K; j++) {
        sum += input_image[(r+i)*(N - K + 1) + (c+j)] *
             kernel[i*K+j];
      }
    }
    output_image[r+dead_rows][c+dead_cols] = (sum /
         normal_factor);
  }
 }
```

Listing 4 shows the result of applying rule R_0 to the code in Listing 3. It can be seen that the array `output_image` is transformed into a one dimensional array and the vector of features associated to code changes accordingly.

Listing 4. Transformation step 3

```
// ABSTR: [3, 0, 0, 0, 0, 0, 1, 0, 0, 1, 1, 0, 2, 4, 0]

int dead_rows = K / 2;
int dead_cols = K / 2;

int normal_factor = K * K;

for(r = 0; r < N - K + 1; r++) {
  for(c = 0; c < N - K + 1; c++) {
    sum = 0;
    for (i = 0; i < K; i++) {
      for (j = 0; j < K; j++) {
        sum += input_image[(r+i)*(N - K + 1) + (c+j)] *
             kernel[i*K+j];
      }
    }
    output_image[(r+dead_rows)*(N-K+1) +
         (c+dead_cols)] = (sum / normal_factor);
  }
 }
```

Listing 5 shows the result of applying rule R_1 to the code in Listing 4. It can be seen that the two outermost loops are collapsed into one for loop but keeping an iteration space with the same number of iterations. Now the code abstraction reflects the change since the number or loops has decreased by one.

Listing 5. Transformation step 4

```
// ABSTR: [2, 0, 0, 0, 0, 0, 1, 1, 0, 1, 1, 0, 2, 4, 0]

int dead_rows = K / 2;
int dead_cols = K / 2;

int normal_factor = K * K;

for (z = 0; z < (N - K + 1)*(N - K + 1); z++) {
    sum = 0;
    for (i = 0; i < K; i++) {
      for (j = 0; j < K; j++) {
        sum += input_image[((z / (N - K +
             1))+i)*(N - K + 1) + ((z % (N - K +
             1))+j)] * kernel[i*K+j];
      }
    }
    output_image[((z / (N - K +
         1))+dead_rows)*(N - K + 1) +
          ((z % (N - K +
               1))+dead_cols)] = (sum / normal_factor);
}
```

Table 1 shows the resulting values of the *state-action* table (Q) for the simple transformation sequence described before. In Table 1 there are as many rows as states obtained from the evaluation of $SM(A(C_i))$ for each code C_i, where C_0 is the initial code and C_4 is the final code classified as ready code for FPGA. Table 1 shows the learned sequence composed of four steps: three consecutive applications of rule R_0 and one application of rule R_1. The Q values of this sequence are highlighted in blue. It can be seen how Q values decreases from the state $SM(A(C_3))$, with the highest reward, down to the initial state $SM(A(C_0))$. This decay behavior in the sequence where Q values decrease as states get further from final states, is caused by the discount factor (γ) introduced in Equation 1. It should be noted that Q values for final states are not updated by the recursive expression in Equation 1. Thus, the final state $SM(A(C_4))$ keeps the initial value of 1.

	$AM(R_0)$	$AM(R_1)$	$RS(C_i)$
$SM(A(C_0))$	17.03718317	16.21544456	R_0
$SM(A(C_1))$	17.25327145	16.80486418	R_0
$SM(A(C_2))$	17.51541052	16.7189079	R_0
$SM(A(C_3))$	16.72942327	17.78007298	R_1
$SM(A(C_4))$	1.	1.	-

Table 1. Values learned for Q table

The example used in this section shows the transformation of a piece of C code. However, the fact that the machine learning methods used work on program abstractions makes the approach generic and suitable for other imperative languages like FORTRAN, which is also widely used in scientific computing. The application of the approach to other languages would require changes to the tool described in Section 4 in order to account for some specific syntactical patterns of a particular programming language. Nevertheless, most of the abstraction features identified and described in Section 4 are also applicable to other imperative languages since they account for common aspects like: control flow, data layout, data dependencies, etc.

6. Results

This section shows some preliminary results obtained for the machine learning -based transformation strategies. We have performed some measurements to verify our claims in Section 5.2 regarding the non-monotonic behavior of non-functional properties for codes obtained from a transformation sequence and how transformed codes with non-optimal values can produce better performant final codes. In order to show this behavior we have identified four different transformation sequences for a use case application performing image compression using the discrete cosine transform.

Each of the identified sequences produces C code which is adequately shaped to be mechanically translated to OpenCL and executed on a GPU. First, we have measured the average execution time of 30 runs for each intermediate code of each sequence. These results are reported in Figure 4 which shows the evolution of the execution times. It can be seen that execution times of sequences do not follow a monotonic behavior and have constant changes. As expected, we also obtain different sequential codes *ready* to be translated to OpenCL and with different execution times.

As a second step, we have measured the performance of the OpenCL version generated from each sequence. Figure 5 shows the execution times of the original sequential program and OpenCL versions. Figure 5 also reports, on top of each bar, the speed up of each OpenCL version with respect to the sequential program, showing that the version obtained from sequence 4 is the fastest one accelerating the computation by a factor of 2.53x. Now, looking at Figures 4 and 5 together it can be seen the absence of correlation between the non-functional property measured (i.e. execution time) of each final sequential code and the performance of the

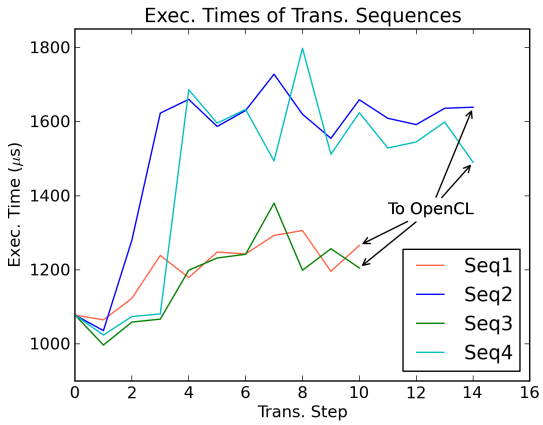

Figure 4. Execution times for transformation sequences

corresponding OpenCL version. In fact, sequence 4 with the second highest execution time produces the fastest GPU implementation. Based on these results it can be concluded that an effective method must be used to discover and learn the uncorrelated relation between final sequential codes and parallel versions for a range of target platforms. We have decided to base our approach on reinforcement learning , since it is driven by final performance measurements rather than on intermediate values which can lead to suboptimal results.

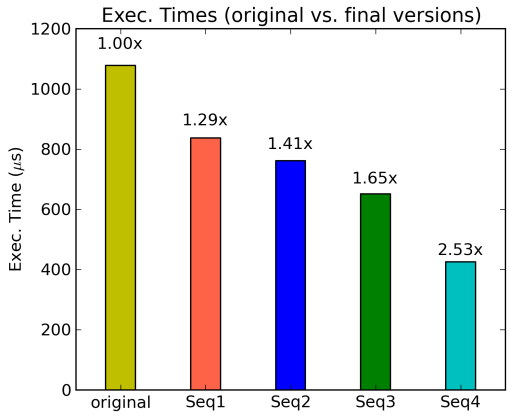

Figure 5. Execution times for OpenCL versions

We have also performed a preliminary evaluation of reinforcement learning as a technique to guide a rule-based program transformation tool. In order to do the aforementioned evaluation we have selected four use case applications and identified different transformation sequences leading to codes that can be mechanically translated to OpenCL and executed on a GPU. These four applications and its corresponding transformation sequences have been used as training set. One of the applications of the train set is the image compression program, mentioned before, denoted as *compress*. Other program of the train set, denoted as *rgbFilter*, is an image filter that splits the different color channels of a given RGB input image in different images. Two more image processing applications complete the train set: one, denoted as *edgeDetect*, detects edges in an image using a Sobel filter and another one, denoted as *threshold*, performs image segmentation over an input image given a threshold value. Once the training set is defined, the transformation engine based on reinforcement learning requires to tune the two parameters that appear in Equation 1, i.e. the learning rate (α) and discount factor (γ). For this purpose, an empirical study was performed in which parameter values leading to transformation sequences providing the fastest OpenCL versions were selected. According to this criteria, a value of 0.5 was used for α and 0.6 for γ. Also, reward values have been chosen in order to give a higher reinforcement to those sequences leading to better performant final codes. In our problem, reward values used for the best sequences of each use case application are substantially higher, with a ratio of 1:100, with respect to the rest of transformation sequences.

After training, three different use case applications were used as predict set. The applications in the predict set were mechanically transformed according to the previously learned sequences and finally translated by hand into OpenCL. Independently, OpenCL versions of the initial C code were written. The evaluation of the the reinforcement learning approach was made by comparing the performance of the hand-coded versions with that of the versions mechanically generated from the transformed codes produced by the sequences learnt by reinforcement learning .

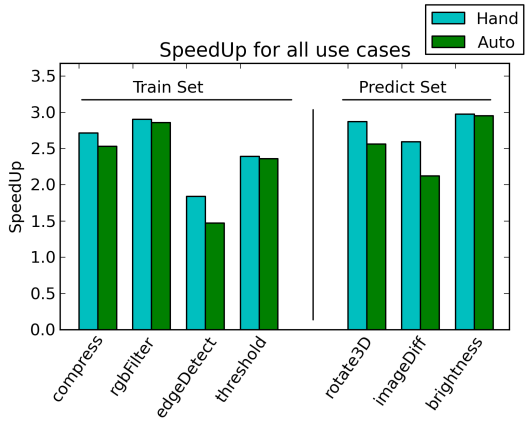

Figure 6. Speed up for train and predict sets

Figure 6 shows the results obtained for both the train and predict sets in terms of speed up factor of the hand and mechanically generated OpenCL versions over the original sequential program. By comparing the results for both the train and the predict set we can assess the response of the transformation system after the learning phase. Looking at results in Figure 6 we can see that the transformation sequences leads to parallel versions that provide comparable acceleration factors with respect to hand coded versions. Although these preliminary evaluation is based on a small sample, it shows that our approach seems promising to tackle the problem of rule-based program transformation systems.

The results discussed in this section show a preliminary evaluation for GPU platforms. However, the same approach can be followed in order to target different platforms which are present in Heterogeneous systems. In this way, a separate *state-action* table can be used for learning adequate transformation sequences for each target platform and rewards would be assigned based on the performance of the final code generated for each platform.

7. Conclusions and Future Work

In this paper we have proposed a machine learning -based approach to learn heuristics for guiding the code transformation process of a rule-based program transformation system (Tamarit et al. 2015, 2016). This type of systems pose a number of problems such as the search-space exploration problem arising from the application of transformation rules in arbitrary orders or the definition of a stop

criteria of the transformation system. For the latter we propose the use of classification trees and for the former we propose a novel approach based on reinforcement learning . We have also performed a preliminary evaluation of the approach, which provided promising results that demonstrate the suitability of the approach for this type of transformation systems.

As a future work we plan to continue expanding the set of use case applications considered for training the different machine learning techniques used. As a consequence, we expect to enrich the code features identified for obtaining program abstractions in order to capture new code patterns. We also expect that the increase in the training set will result in better prediction outcomes, but it will also increase the complexity of efficiently using the learning techniques. In the case of reinforcement learning having more learning cases results in a bigger state space. For that reason we plan to use common methods to reduce the number of states like clustering or principal component analysis techniques. We also plan to explore another features of reinforcement learning like the use of different learning rates for different states or transformation sequences in order to learn and converge faster towards transformed codes providing better performant final versions.

We also plan to define *multi-objective* reinforcement learning rewards based on different non-functional properties like energy consumption, resource usage, etc. and even a combination of multiple of these properties. This future line would permit the definition of transformation strategies that try to optimize different aspects and generate, for example, a final code among the fastest ones that consumes the least amount of energy. Also, the introduction of weights into the *multi-objective* rewards would offer programmers the flexibility to select which non-functional property or set of properties they want to focus on for generating the final code.

Acknowledgments

Work partially funded by EU FP7-ICT-2013.3.4 project 610686 POLCA, Comunidad de Madrid project S2013/ICE-2731 N-Greens Software, and MINECO Projects TIN2012-39391-C04-03 / TIN2012-39391-C04-04 (StrongSoft), TIN2013-44742-C4-1-R (CAVI-ROSE), and TIN2015-67522-C3-1-R (TRACES).

References

F. Agakov, E. Bonilla, J. Cavazos, B. Franke, G. Fursin, M. F. P. O'Boyle, J. Thomson, M. Toussaint, and C. K. I. Williams. Using machine learning to focus iterative optimization. In *Proceedings of the International Symposium on Code Generation and Optimization*, CGO '06, pages 295–305, Washington, DC, USA, 2006. IEEE Computer Society. ISBN 0-7695-2499-0. doi: 10.1109/CGO.2006.37. URL http://dx.doi.org/10.1109/CGO.2006.37.

O. S. Bagge, K. T. Kalleberg, E. Visser, and M. Haveraaen. Design of the CodeBoost Transformation System for Domain-Specific Optimisation of C++ Programs. In *Third International Workshop on Source Code Analysis and Manipulation (SCAM 2003*, pages 65–75. IEEE, 2003.

A. Brown, W. Luk, and P. Kelly. Optimising Transformations for Hardware Compilation. Technical report, Imperial College London, 2005.

E. Di Napoli, D. Fabregat-Traver, G. Quintana-Orti, and P. Bientinesi. Towards an efficient use of the BLAS library for multilinear tensor contractions. *Applied Mathematics and Computation*, 235:454–468, May 2014.

D. Fabregat-Traver and P. Bientinesi. Application-tailored linear algebra algorithms: A search-based approach. *International Journal of High Performance Computing Applications (IJHPCA)*, 27(4):425 – 438, Nov. 2013.

F. Franchetti, Y. Voronenko, and M. Püschel. FFT Program Generation for Shared Memory: SMP and Multicore. In *Supercomputing (SC)*, 2006.

L. P. Kaelbling, M. L. Littman, and A. P. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.

J. G. Koomey. Worldwide Electricity Used in Data Centers. *Environmental Research Letters*, 3(3):034008, 2008. URL http://stacks.iop.org/1748-9326/3/i=3/a=034008.

G. Mariani, G. Palermo, R. Meeuws, V. M. Sima, C. Silvano, and K. Bertels. Druid: Designing reconfigurable architectures with decision-making support. In *19th Asia and South Pacific Design Automation Conference, ASP-DAC 2014, Singapore, January 20-23, 2014*, pages 213–218, 2014. doi: 10.1109/ASPDAC.2014.6742892. URL http://dx.doi.org/10.1109/ASPDAC.2014.6742892.

S. Marsland. *Machine Learning: An Algorithmic Perspective*. Chapman & Hall/CRC, 1st edition, 2009. ISBN 1420067184, 9781420067187.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

G. Pekhimenko and A. Brown. Efficient program compilation through machine learning techniques. In K. Naono, K. Teranishi, J. Cavazos, and R. Suda, editors, *Software Automatic Tuning*, pages 335–351. Springer New York, 2010. ISBN 978-1-4419-6934-7. doi: 10.1007/978-1-4419-6935-4_19. URL http://dx.doi.org/10.1007/978-1-4419-6935-4_19.

T. Schaul, J. Bayer, D. Wierstra, Y. Sun, M. Felder, F. Sehnke, T. Rückstieß, and J. Schmidhuber. PyBrain. *Journal of Machine Learning Research*, 2010.

S. Schupp, D. Gregor, D. Musser, and S.-M. Liu. Semantic and behavioral library transformations. *Information and Software Technology*, 44(13): 797–810, 2002.

R. M. Stallman and GCC Developer Community. *Using The Gnu Compiler Collection: A Gnu Manual For Gcc Version 4.3.3*. CreateSpace, Paramount, CA, 2009. ISBN 144141276X, 9781441412768.

J. E. Stone, D. Gohara, and G. Shi. OpenCL: A Parallel Programming Standard for Heterogeneous Computing Systems. *IEEE Des. Test*, 12 (3):66–73, May 2010. ISSN 0740-7475. doi: 10.1109/MCSE.2010.69. URL http://dx.doi.org/10.1109/MCSE.2010.69.

S. Tamarit, G. Vigueras, M. Carro, and J. Mariño. A Haskell Implementation of a Rule-Based Program Transformation for C Programs. In E. Pontelli and T. C. Son, editors, *International Symposium on Practical Aspects of Declarative Languages*, number 9131 in LNCS, page TBD. Springer-Verlag, June 2015.

S. Tamarit, J. Mariño, G. Vigueras, and M. Carro. Towards a Semantics-Aware Code Transformation Toolchain for Heterogeneous Systems. In *First Workshop on Program Transformation for Programmability in Heterogeneous Architectures (PROHA'16)*, March 2016.